The Role of Uninformed Individuals in Making the Right Group Decisions

Carissa Flocken, University of Michigan
Ted Carmichael and Mirsad Hadzikadic, University of North Carolina at Charlotte

Abstract. Is being uninformed always bad for the group decision making? Under what circumstances can some degree of ignorance be in fact positive in the aggregate? What if both quality and efficacy of the decision making process matters? This papers attempts to answer these questions. In addition, the simulation and modeling platform developed for this research was also used to verify results of similar experiments performed in real life.

1 Introduction

This paper discusses the role of individual knowledge in terms of the collective and individual quality of a solution. In lieu of the independence of knowledge requirement for the wisdom of the crowds model [12], we explore a possible range of knowledge that an individual may have of a given subject. We ask what role does a lack of information or knowledge when an individual is polled for a decision in an animal or human group?

Prior research outlines a narrative of groups containing individuals who are uninformed about a decision, such as movement targets, hunting strategies, nest sites, or migration routes [2,16,10,14,15,5], and are vulnerable to manipulation by smaller, informed, opinionated soles or factions within the group [1,12,13,8]. This belief proposes that uniformed or naive "individuals destabilize the capacity for collective intelligence in groups by allowing or facilitating the propagation of extremist opinions in populations" [8]. However, Ian D. Couzin and his fellow researchers [8] found that uniformed individuals play a central role in achieving democratic consensus. They took a theoretical approach to this position and found that in all of their models, "an entrenched minority is capable of exerting substantial influence by biasing the perceived consensus. Because they exhibit intransigence or intrinsic bias, however, uninformed individuals will lend support to, and tend to amplify, a numerical advantage (even a slight one). If sufficiently numerous, they reduce the effect of intransience and inhibit the capacity for the minority to take hold, thus returning control to the numerical majority. Consequently, even a small change in the number of uninformed individuals can dramatically alter the outcome of consensus decisions." They emphasizes that this process will tend to impede any strong minority preference, "regardless of the intrinsic quality or value of that view" by enforcing equal representation and promoting a democratic outcome [8].

2 Experimentation

2.1 Blue & Gold World

In order to evaluate influence of uninformed individuals on the optimal group decision making, we created the Blue & Gold world, a simulation of individual agents/decision makers in the NetLogo environment. Blue & Gold is a 10 x 10 torus of 100 agents represented by cells on a grid. Each agent is initialized as either Blue or Gold. In the first condition these agents are initialized at random; in the second, the user determines the number of each color. The experiment is run until a consensus is reached; i.e., all agents are Blue or all are Gold. Each agent has an opportunity to change color based on a poll of some number of other agents selected at random. There are two ways to measure the "cost" of such polling. The first is the absolute number of agents that must be polled before consensus is reached. The second is the number of polling rounds that are needed.

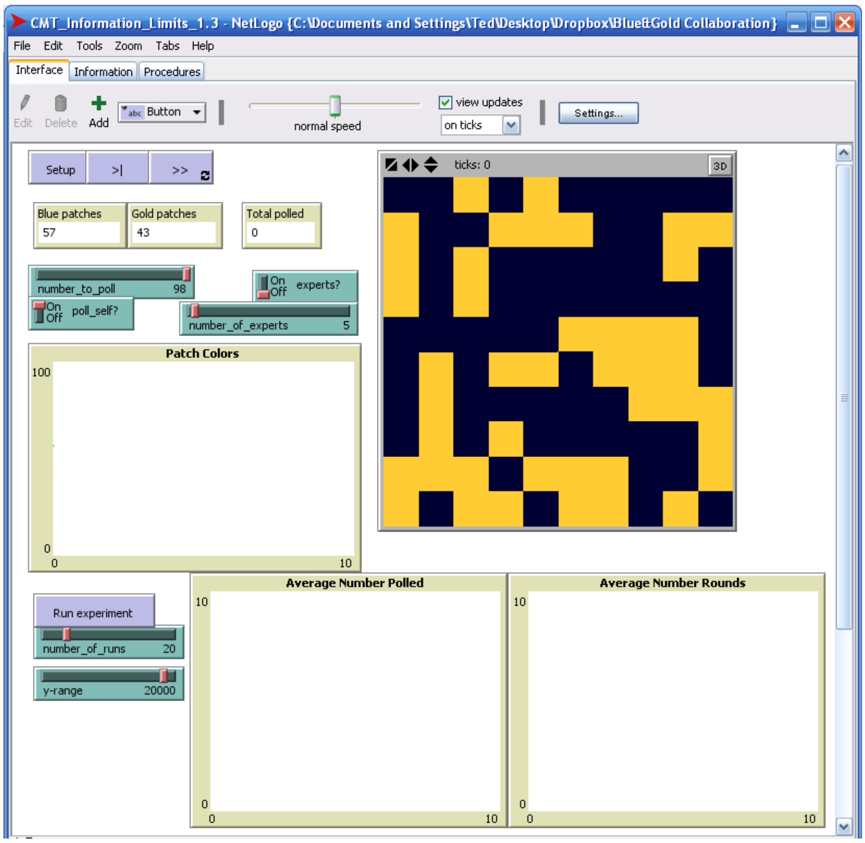

Figure 1 Screen shot of the Blue and Gold world model interface.

During a single round, every agent has the opportunity to poll a set number of other agents. These agents are selected at random. Thus, when the number of agents to be polled is set, for example, at 10, each agent will tally how many of those polled are blue and how many are gold. Then this agent adds its own color to the totals, and will change its color to that of the simple majority. In the case of a tie, the agent does not change its color. During a single round of voting, all tallies are done first, and then all necessary color changes are done based on these tallies. Thus, the polling process in each of these influence steps are completed in parallel, rather than serially.

An experiment is run by incrementing the number polled by each agent during each round of voting, from 2 to 98. At each level of number polled, 40 repeated runs are conducted and the results are averaged for that level, in terms of both the total number polled to reach consensus, and the number of rounds needed.

Experiment 1: all agents are initialized to either Blue or Gold at random; 100 runs each, from 2 to 98 agents polled by each agent, each round.

The number of rounds needed to reach a consensus is high when the number of agents polled is very low, and drops off quickly as the size of the group polled is increased. The resulting curve is very close to exponential, as seen below.

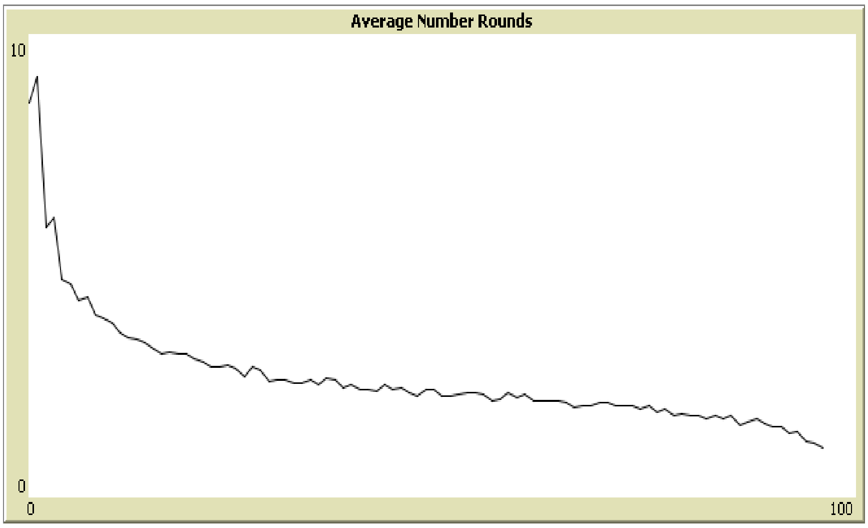

Figure 2 Number of rounds needed to reach a consensus, for each set of parameters, averaged across 50 runs for each setting.

The curve starts at 8.53 rounds when only 2 other agents are polled, and ends at a low of 1.05 rounds when 98 other agents are polled. (The number of rounds needed when 3 other agents are polled is slightly higher than with 2; this is because ties are possible when the number polled, plus one "poll" for the agent's own color, is even.) The Y-axis is the number of rounds and the X-axis is the number polled by each agent during each round.

In terms of the total number of agents polled, the results are somewhat different. As can be seen in the graph below, the most efficient is when each agent polls only two other agents each round. The average of the total number polled, across 100 runs, is 1706. This total slowly increases, as the number polled each round increases, for most of the experiments, maximizing around ~15,000. Near the end, however, as the number polled nears the maximum of 98, the total number polled begins to decline. At exactly 98, a total (average) of 10,290 agents are polled. This is due to the fact that very few rounds – usually just one – are needed to reach consensus.

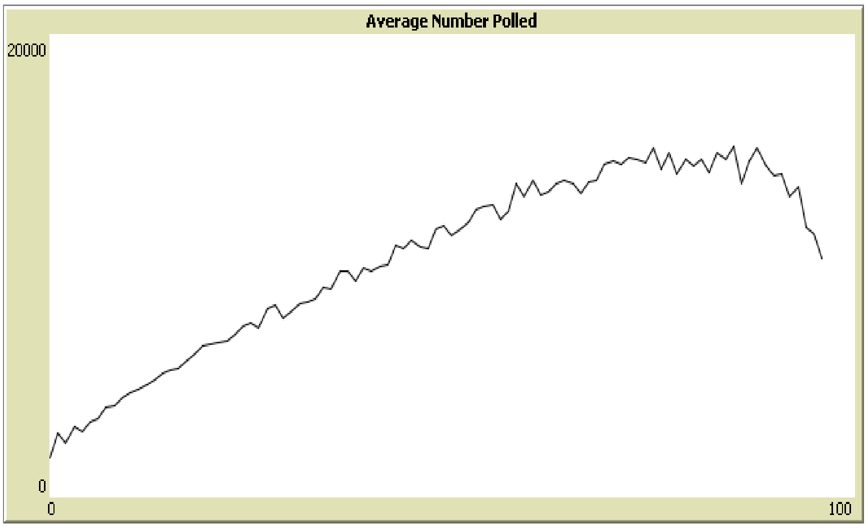

Figure 3 Total number of polls needed to reach a consensus, for each set of parameters, averaged across 50 runs for each setting.

Experiment 2: Exactly 50 of each color at initialization.

When the initialized number of each color is not random, but rather exactly half of the agents are Blue and half Gold, the results are very similar to experiment 1. Across 100 runs of each setting, the lowest average number of rounds is 9.14, and the lowest average number polled is 1828. The maximums, however, are larger. This is because in the random initialization, sometimes Blue or Gold already has more than half of the total number of agents; thus, reaching a consensus is often both easier and less costly.

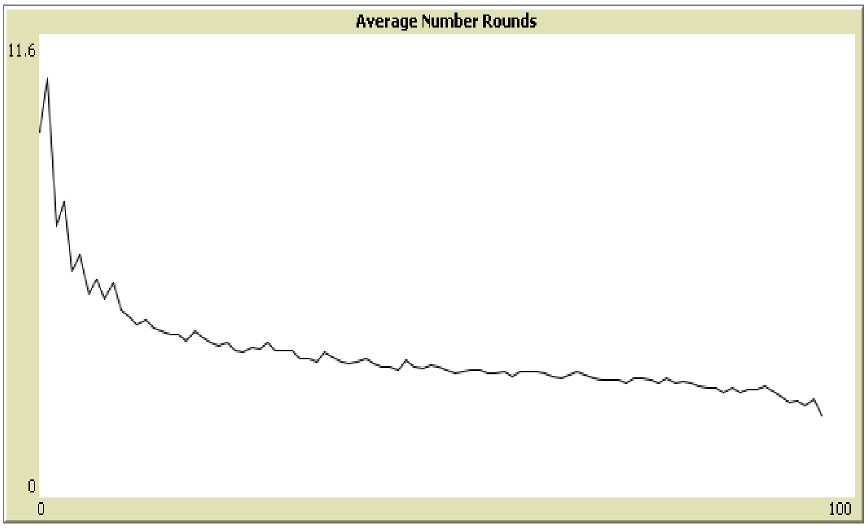

Figure 4 Number of rounds needed to reach a consensus, for each set of parameters, averaged across 50 runs for each setting.

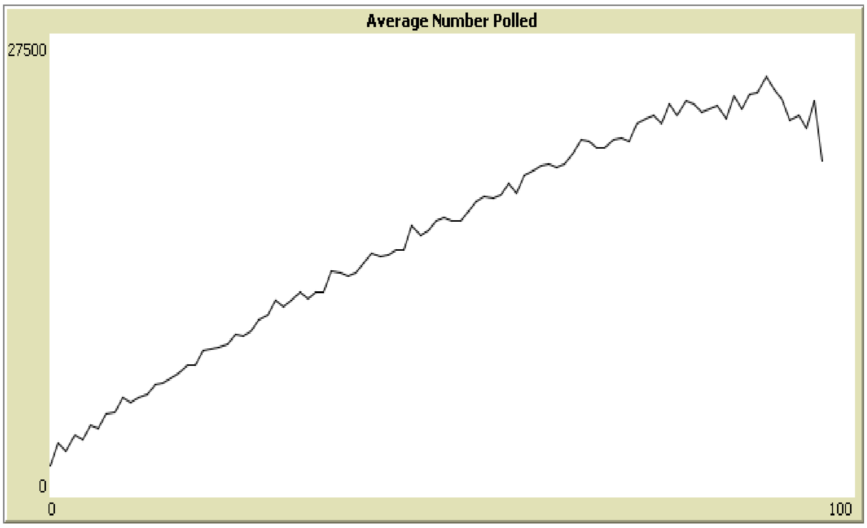

Figure 5 Total number of polls needed to reach a consensus, for each set of parameters, averaged across 50 runs for each setting.

These experiments demonstrate that limited knowledge of agents, expressed in their lack of awareness of other agents' "opinions," may be useful for determining the opinion of the majority in a cost effective way. In this particular case the cost was expressed in terms of either: the number of comparisons with other agents in order to determine their opinion/color; or, the number of rounds needed for these comparisons, in order to reach a consensus.

2.2 Couzin et al's Model

After completing this set of experiments it occurred to the authors that the same Blue & Gold environment could be used to replicate and expand Couzin et al's fish model. In this model it holds that N1 > N2; where N1 is the majority position population (they have a weak preference) and N2 is the minority position population (they have a strong preference). In Couzin's experiment, N1 = 6 and N2 = 5.

The Couzin, et. al, model uses two radii for local influence, such that agents avoid others that are too close, and moves closer in alignment to those further away. The direction is continuous, such that it can range between the blue target and the gold target. It is the movement towards alignment in the second radius that produces a consensus in the direction of travel, towards one of the two targets.

Our model instead uses a discrete choice between the two targets. During each step, each agent picks one other agent at random and, based on a random-float between [0, 1) as compared to its own 'susceptibility,' will change (or not) its direction to match the selected agent. Each agent then moves forward towards one of the two targets. This process repeats until all the agents face the same target. If no consensus is reached before all the fish reach one of the two targets, then the result is a "split decision."

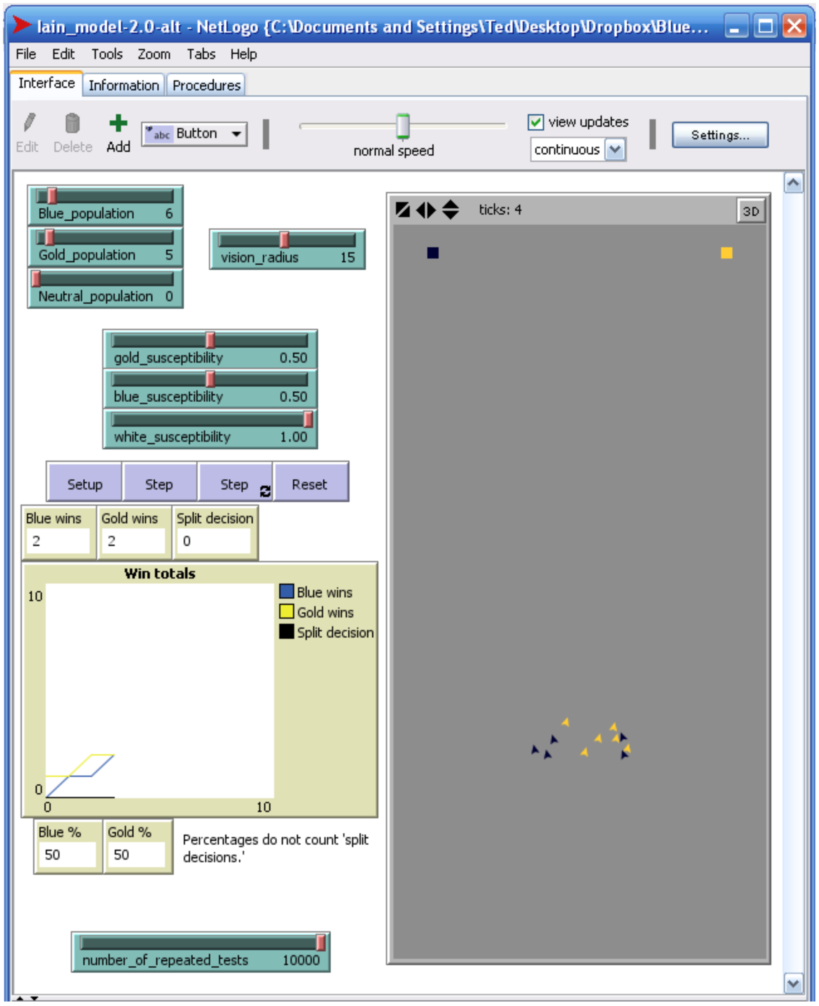

Figure 6 Screenshot of the model interface.

When "uninformed" agents are added, these agents cannot (initially) influence other agents. Their susceptibility is set to the maximum (= 1.0), and thus they will always change to match the direction of the pair-wise counterpart. They can, however, influence other agents (to the degree allowed by those agents susceptibility) once they have adopted one of the two directions (i.e., towards the blue target or the gold target).

The following experiments were conducted with our model:

Baseline:
Blue population = 6; Gold population = 5; "susceptibility" same for both populations; 10,000 runs for each setting group.

Susceptibility	Result
0.30	Blue wins 54.59%; Gold wins 45.41%; (split = 1.88%)
0.40	Blue wins 54.1%; Gold wins 45.9%; (split = 0.22%)
0.50	Blue wins 54.46%; Gold wins 45.54%; (split = 0.02%)
0.60	Blue wins 55.42%; Gold wins 44.58%; (split = 0)
0.70	Blue wins 54.39%; Gold wins 45.61%; (split = 0)
0.80	Blue wins 54.67%; Gold wins 45.33%; (split = 0)
0.90	Blue wins 54.65%; Gold wins 45.35%; (split = 0)

In this baseline condition there is a small, but distinct, advantage for the larger population in the Blue group, across matching susceptibility rates from 0.30 and above. At a susceptibility rate of 0.10 the number of "split decisions" (no consensus by the time the fish have traversed the entire distance) was high, and therefore susceptibility below 0.30 was not explored further. The majority advantage was largely consistent across all subsequent levels of susceptibility, from 0.30 to 0.90.

1st condition: Blue group is MORE susceptible to influence, and the gold group is LESS susceptible to influence. 10,000 runs for each condition. (The number of rounds where no consensus was reached was never higher that 0.45% of the total; thus, these numbers are excluded from the next two tables.)

Susceptibility	Result
Blue = 0.55; Gold = 0.50	Blue wins 51.2%; Gold wins 48.8%
Blue = 0.60; Gold = 0.50	Blue wins 48.12%; Gold wins 51.88%
Blue = 0.65; Gold = 0.45	Blue wins 42.91%; Gold wins 57.09%
Blue = 0.70; Gold = 0.40	Blue wins 37.38%; Gold wins 62.62%
Blue = 0.75; Gold = 0.35	Blue wins 31.08%; Gold wins 68.92%
Blue = 0.80; Gold = 0.30	Blue wins 24.98%; Gold wins 75.02%

The Gold group consistently prevailed more often than the Blue group, immediately offsetting the Blue group majority advantage, except in the condition with the lowest susceptibility difference between the two groups (0.55 for Blue and 0.50 for Gold). As the Gold group advantage becomes more pronounced – i.e., as their conviction becomes much stronger than that of the Blue group – the results skew more in their favor.

This result qualitatively matches the Couzin, et. al, results, even though a different ABM method, with a different influence mechanism, is used to test the influence of the more-committed minority.

2nd condition: same as 1st condition, but with a population of "uninformed" and highly susceptible agents (N3 = 10; susceptibility = 1.00). 10,000 runs for each setting.

Susceptibility	**Result**
Blue = 0.55; Gold = 0.50; White = 1.00	Blue wins 52.92%; Gold wins 47.08%
Blue = 0.60; Gold = 0.50; White = 1.00	Blue wins 50.11%; Gold wins 49.89%
Blue = 0.65; Gold = 0.45; White = 1.00	Blue wins 45.18%; Gold wins 54.75%
Blue = 0.70; Gold = 0.40; White = 1.00	Blue wins 40.67%; Gold wins 59.33%
Blue = 0.75; Gold = 0.35; White = 1.00	Blue wins 37.28%; Gold wins 62.72%
Blue = 0.80; Gold = 0.30; White = 1.00	Blue wins 32.05%; Gold wins 67.95%

The Gold group will still prevail when the difference in susceptibility between the two groups is large. However, the Blue group will prevail when this difference is small. Further, the effect of adding a highly susceptible and uninformed group will reduce the susceptibility advantage of the Gold group under all conditions.

This result matches the Couzin, et. al, results in that an uninformed minority will influence the direction of the overall consensus towards the initial opinion that has a plurality. This effect is not enough to overcome the intransigence of the minority opinion when the difference in susceptibility is large between the Blue group and the Gold group. However, the direction of this effect is qualitatively the same as the Couzin, et. al, results, in terms of lessening the influence of the minority. This is true even though a different method of influence, and a different ABM mechanism, is used.

3 Conclusion

Agent-based modeling offers a plausible and rich environment for evaluating decision-making strategies and their consequences. The research presented in this paper evaluates the role of uninformed individuals in helping the group reach the consensus decision in the most effective way. Blue & Gold presents a simple yet powerful environment for both evaluating various models of decision making and replicating existing real-life experiments. Even more importantly, this environment allows researchers to conduct experiments that are either too costly or take too long a time to evaluate in their natural settings.

References

[1] J. M. Buchanan, G. Tullock, The Calculus of Consent: Logical Foundations of Constitutional Democracy (Liberty Fund, Indianapolis, IN, 1958).
[2] Camazine, S., Deneubourg, J. L., Franks, N. R., Sneyd, J., Theraulaz, G. & Bonabeau, E. 2003 Self-organization in biological systems. Princeton, NJ: Princeton University Press.

[3] Conradt, L. 1998 Could asynchrony in activity between the sexes cause intersexual social segregation in ruminants? Proc. R. Soc. B 265, 1359–1363. (doi:10.1098/rspb.1998. 0442)
[4] Conradt, L. & Roper, T. J. 2003 Group decision-making in animals. Nature 421, 155–158. (doi:10.1038/nature 01294)
[5] Conradt, L. & Roper, T. J. 2005 Consensus decision making in animals. Trends Ecol. Evol. 20, 449–456. (doi:10.1016/ j.tree.2005.05.008)
[6] Conradt, L. & Roper, T. J. 2007 Democracy in animals: the evolution of shared group decisions. Proc. R. Soc. B 274, 2317–2326. (doi:10.1098/rspb.2007.0186)
[7] Conradt, L. & Roper, T. J. 2009 Conflicts of interest and the evolution of decision sharing.
[8] Conradt, J. Krause, I. D. Couzin, T. J. Roper, Am. Nat.173, 304 (2009).
[9] Couzin, I. D., Krause, J., Franks, N. R. & Levin, S. A. 2005 Effective leadership and decision-making in animal groups on the move. Nature 433, 513–516. (doi:10.1038/ nature03236)
[10] Couzin, I. et Al. Uniformed Individuals Promote Democratic Conseenus In Animal Groups. Science 334 (2011).Phil. Trans. R. Soc. B 364, 807 –819.(doi:10.1098/rstb.2008.0257)
[11] C. MacKay, 1841, "Extraordinary Popular Delusion and the Madness of Crowds", ISBN 1853263494, 9781853263491
[12] M. Olson Jr., The Logic of Collective Action: Public Goods and the Theory of Groups (Harvard Univ. Press, Cambridge, MA, 1971).
[13] W. Riker, Liberalism Against Populism: A Confrontation Between the Theory of Democracy and the Theory of Social Choice (W. H. Freeman, San Francisco, 1982).
[14]Seeley, T. D. & Buhrman, S. C. 1999 Group decision making in swarms of honey bees. Behav. Ecol. Sociobiol. 45, 19–31. (doi:10.1007/s002650050536)
[15] Simons, A. M. 2004 Many wrongs: the advantage of group navigation. Trends Ecol. Evol. 19, 453–455. (doi:10.1016/ j.tree.2004.07.001)
[16] Trillmich, J., Fichtel, C. & Kappeler, P. M. 2004 Coordi- nation of group movements in wild Verreaux's sifakas (Propithecus verreauxi). Behaviour 141,1103–1120. (doi:10.1163/1568539042664579)

ODD protocol – Blue & Gold Model (model #1)

1. Purpose.

To explore consensus building with a simple network of agents. This model instantiates each agent with one of two colors, and then allows multiple rounds of "polling" other members of the network. The agents' goal is to form a consensus; i.e., reach the state where all agents are the same color. They do this by polling a sub-set of other members and using a simple majority of those polled to determine their color for the current round of polling. The use can choose the size of the sub-set to be polled. Multiple rounds of polling are conducted until a consensus is reached.

2. Entities, state variables, and scales.

The model has 100 agents (patches, in NetLogo terminology) in a 10 x 10 grid. Each patch has one of two states: Blue or Gold. During each round of polling, each agent creates a sub-set of other patches (chosen at random) and queries these patches as to which state they are in. Some agents, as designated by the user, can be "expert" agents. These will poll ALL other agents during each round, and base their color on this information.

3. Process overview and scheduling.

In one version of the model, each agent's state (color) is chosen at random between Blue and Gold. In a second version, the number of Blue agents (and hence, the number of Gold agents also) is chosen by the user. Even when the number is chosen by the user, the distribution of initial Blue and Gold agents is random.

A "round" of polling is defined as: every agent chooses a sub-set of other agents to poll. The size of this sub-set is determined by the user, with the slider labeled *number_to_poll*, and is the same size for all agents (except "experts," if the switch "experts?" is turned ON). The polling is done in parallel, in that all agents conduct their poll first, and then all the agents change their state (their color) based on the results of their individual polls.

The agents use a simple majority among those polled to determine their next state. They include their own color in these tallies, unless *poll_self?* is turned OFF. If experts? is turned ON, then the slider *number_of_experts* determines how many agents are experts. These experts are chosen at random at the beginning of each experiment from the pool of 100 agents. An agent will change its state only if a simple majority of those polled have a different state than that agent. A tie will result in that agent staying with its current state.

Additional rounds of polling will occur until every agent has the same state, and thus a consensus is reached. The user can also control the simulation such that a button is pushed for each round of polling.

4. Design concepts.
The basic principle of this model is to test how quickly a consensus can develop as the agents poll other agents as to their current state, and change their own state based on the results of these polls.

Emergence: the consensus among all the agents emerges based on the interaction of these agents as represented by the polling process.
Adaptation: the agents adapt to the majority of those polled by collecting data on other agents and letting the majority of those polled decide its next state.
Objectives, *Learning*, and *Prediction* are not used in this model.
Sensing: sensing is captured by the polls that each agent conducts. Thus, each agent can 'sense' the current state of the other agents that it polls.
Stochasticity: this is captured by: the initial state of each agent is random. (This also includes the number of each color, unless the exact number of Blue and Gold agents is determined by the user. In the first case, the number will tend to be 50 of each, but will follow a normal distribution around a mean of 50. In the second case, only the distribution of each color is random.) The sub-set of other agents to be polled each round is also selected randomly. There is no preference based on locality in choosing this sub-set. If experts are used, these are selected randomly from the pool of 100 agents.
Collectives are not used in this model.

Observation: there are two main plots. The user selects how many experiments will be conducted for each size of sub-sets to be polled. This only applies if the user clicks the *Run experiment* button, which increments through every size sub-set from 2 through 98. If number_of_runs is, for example, 20, then 20 experiments will be conducted with each agent polling two other agents each round. Then 20 more experiments will be conducted for three other agents polled, and four, and so on. The first plot averages the total number of agents polled, accumulated across all rounds, across all 20 runs. This represents one data point on this plot. The next data point will average the number polled for the next size sub-set across 20 runs, and so on. The second plot records the average number or rounds needed to reach consensus, across 20 runs, for each size sub-set.
Explanation: As each agent can choose to poll *any* other agent, then all agents are therefore in the neighborhood of that agent. Thus, this model environment represents a fully connected graph. Other design choices were considered. However, if the pool of potential neighbors is based on locality, then it is possible that a full consensus is never reached. This phenomena is well known in the literature where it discusses the persistence of minority opinion. When this happens in our model, however, the simulation would never reach the end state. We decided to exclude this possibility by using a fully connected graph. For the same reason, the size of sub-sets to be polled can never be 99, which, of course, would represent ALL of the other agents, and could also lead to an infinite loop.

5. Initialization.
Each agents is assigned a color, either Blue or Gold. All other pertinent questions about initialization have been answered in other sections.

6. Input data.
Not used.

7. Submodels
Not used, or described elsewhere.

ODD protocol – Couzin, et. al, replication model (model #2)

1. Purpose.

To explore consensus building with a simple collection of agents. Each agent represents a fish swimming towards a food source. Some fish are attracted to the blue colored target and some are attracted to the gold target. The operator can also add fish that have no preference at all. The fish can influence each other, and do so until a consensus is reached; then all the fish swim towards one target or the other.

2. Entities, state variables, and scales.

The agents in this model represent fish. Each fish has a heading and a color that represents that heading. Scale doesn't matter. There are also two targets: a blue one and a gold one. These two targets represent the two possible headings – and hence, the two possible colors – that each fish can have. Each fish is also given a 'susceptibility' based on which group they are in. This variable is the same within groups, and does not change even as the heading/color can change. "White" fish can also be added at initialization. These fish capture a color/heading from other fish based on their susceptibility. The slider "vision_radius" determines how far away one fish can chose another for the pair-wise influence test. The environment is a rectangle grid of 33 x 65 cells. The x-axis ranges from -16 to +16, and the y-axis from -32 to +32.

3. Process overview and scheduling.

Fish are initialized on one side of the long end of the environment and the targets at the other. During each simulation 'tick' each fish: 1) picks one other fish within its radius; and 2) checks the color/heading state of that fish. If the color/heading is different, then 3) a random number between [0, 1) is generated and compared to the first fish's 'susceptibility' setting. If the random number is less than its susceptibility then it will adopt the color/heading of the selected fish. After this process the fish will move a random distance towards the target that is between 0.8 and 1.3 cell lengths.

This process occurs for each fish, one by one and at random, until all the fish have had an opportunity to be influenced by another and move forward. At the end of each step a stopping condition is checked – if either all the fish have the same target, or all the fish have reached a target, then the simulation ends.

4. Design concepts.

This model was designed to test the concepts described in Couzin, et. al, 2011, regarding the dynamics of influence among three distinct groups: a majority (or plurality) group that has a weak preference; a minority group that has a strong preference; and the effect that a third group with no preference (the "uninformed") can have, in terms of the entire population reaching consensus. Thus, the experiments described in our report match those from Couzin, et. al, 2011 in terms of the size of each population. However, the mechanisms for influence and movement among the fish is different than in that paper.

Emergence: the consensus among all the fish emerges based on the interaction of these fish as they travel across the environment towards one of two targets.
Adaptation: The fish can change their state in terms of their heading, based on influence from other fish.
Objectives, *Learning*, and *Prediction* are not used in this model. Ostensibly, each fish has the objective of reaching the target; however, which of the two targets is reached is of no importance; only the consensus of the group – if one is reached – is of concern, and the targets are used only to differentiate between two choices among the fish.
Stochasticity: The initial position of each fish is random, within a certain area. The pair-wise choice for influence is also chosen randomly, within a set radius from the acting fish. Movement towards the targets is a random distance, within a set range.
Collectives are not used in this model.
Observation: The main measurement in this model is "win totals;" i.e., which color, blue or gold, prevails. This is because not only is consensus important, but which color becomes the consensus. In the original study the blue leaning population is larger than the gold leaning one, yet the gold group has a stronger preference. And so the purpose is to see with view prevails under various conditions.
Explanation: see above.

5. Initialization.
The fish are initialized on one end of the environment, randomly between -5 and +5 on the x-axis, and between -30 and -26 on the y-axis. The two targets are located at cell (-13 30) and cell (13 30). Initial population for each of the three groups can be set to anything, but all experiments reported here have the Blue group initially size 6, the Gold group size 5, and – when included – the White group (no preference) are size 10. White susceptibility is set to reflect "no preference" and is thus set to 1.00.

6. Input data.
Not used.

7. Submodels
Not used, or described elsewhere.